MATHematics illuminated

PROGRAM: 9

Game Theory

Producer: Sean Hutchinson
Host: Dan Rockmore

Produced by Oregon Public Broadcasting for Annenberg Media

Time Code	Audio
00:00	OPENING CREDITS
00:40	HOST We've all heard it said that life is like a game.
00:43	HOST Most games, whether we work in teams or work alone, have well defined rules, with clear benefits for winning and costs for losing. And that makes them something we can think about logically and mathematically.
00:57	HOST But what about life? Can mathematics tell us anything about the competitions and collaborations that happen every day between individuals, groups, nations...even between animals or microbes? From the social sciences to biology, robotics and beyond, the answer is yes. Welcome to Game Theory.
01:29	INSPECTOR So Mr. Blue, we got you dead to rights. Picked up you and Mr. White not a half a block from the scene of the robbery.
01:36	BLUE We were out buying groceries.
01:38	WHITE We were out buying groceries
01:41	INSPECTOR That where you bought this little item?
01:43	BLUE That doesn't prove a thing
01:45	WHITE Doesn't prove anything...
01:48	INSPECTOR Really, what do you think your friend Blue will say about that?
01:52	BLUE/WHITE He won't talk. He better not… INSPECTOR Look I'm gonna lay it out for you… You talk and we let you go. BLUE/WHITE No jail time? INSPECTOR

	Nada. Zip
02:01	BLUE/WHITE What happens to White/Blue? INSPECTOR He gets 90 days. BLUE/WHITE What if he talks, and I don't?
02:06	INSPECTOR Well then he walks and you get 90 days. BLUE/WHITE What if he rats on me and I rat him…? INSPECTOR You both get 60 days. BLUE/WHITE What if neither one of us talks? INSPECTOR Well then it's a light sentence; you both do 30 days. But you need to ask yourself: how much do you trust your buddy?
02:25	BLUE/WHITE Okay. He did it.
02:32	HOST Now that wasn't such a good strategy. Or was it? Both Mr. Blue and Mr. White end up in jail. But with the right combination, one or the other could have been free. Then again, if they had cooperated with each other and kept quiet, they'd still both go to jail, but with an easier sentence. So what's their best strategy? Or is there one?
02:54	**HOST (V.O.)** **Our two criminals are, in fact, caught in what's called The Prisoner's Dilemma - a classic scenario of modern game theory, which came into its own as a part of mathematics in the 20th century.**
03:06	HOST You see, the point is that interactions are strategic, say cooperative or competitive, and how well we do in any given strategy almost always depends on the actions of others. The value of an interaction can be expressed in terms of a cost and a benefit as in the loss or capture of

	piece in a chess game. The great surprise of game theory is that it not only applies to "games" but to interactions in the real world, like in the dilemma facing Mr. Blue and Mr. White.
03:38	HOST To do that, let's take a look at the game these kids are playing.
03:45	HOST It's called Odd-Even, sort of a simple version of rock-paper scissors. One kid takes odd and the other takes even. For each round, the kids choose to reveal either one finger or two. When they add up the number of fingers, if that number's odd, the kid who chose Odd wins all the points. If it turns up even, the kid who chose Even gets all the points.
04:09	HOST In every round, one kid wins, and one kid loses. Pretty simple, and it doesn't seem like there's much strategy going on. But let's look further... The best way to understand what the Odd-Even game looks like in terms of who wins and who loses is to build a grid and look at how each single round, or game, could go.
04:30	**HOST (V.O.)** **Let's put Odd on the left, and Even at the top. So if, the first, Odd chooses 1 and Even chooses 1, Even gets the two points, and we can say, theoretically, that Odd loses two points. We write it like this, starting with Odd's score being--**
04:46	**HOST (V.O.)** **-- minus 2, and Even's score being 2: Even's payoff. The second time, maybe Odd chooses 1 again, and Even chooses 2. Now we've got 3 - an odd number, so Odd gets the points.** **Odd's payoff is 3, Even's cost is minus 3.**
05:06	**HOST (V.O.)** **Third time, Odd chooses 2, Even chooses 1. Odd wins again, and again Odd's payoff is 3, Even's cost is minus 3. Fourth time they choose 2 and 2 and so on. Even wins.**
05:18	HOST Now, if we're trying to decide on a best strategy, we actually have to do a little algebra and figure out the probability of each solution turning up.
05:28	HOST Now here's where we see the magic of math. It turns out that if odd plays one 7/12ths of the time, odd will actually accumulate more points

	over time—winning the game. This is an example of a “mixed strategy” because Odd has to mix up what he does.
05:45	**HOST (V.O.)** **In fact, if you do only one thing all the time, your odds of winning aren’t going to increase.** **Just the opposite in the long run, because your opponent’s going to figure out pretty quickly what you’re doing.**
06:00	HOST This kind of payoff matrix does help us see that our instinct for not making the same choice all the time is also a mathematically sound one.
06:08	HOST Odd-Even is an example of what we call a zero-sum game -- “I win you lose” -- a player benefits only at the expense of others. If you add the payoff and benefit for each hand, they add up to zero.
06:22	**HOST (V.O.)** **But most games are non-zero-sum: a gain by one player doesn’t necessarily mean a loss by another player -- as in Blue and White’s Prisoner’s Dilemma. Let’s take a look at their payoff matrix to see if there’s a best strategy for their non-zero-sum game:**
06:38	**HOST (V.O.)** **C stands for Cooperate - the choice to keep quiet. D stands for Defect - the choice to rat the other person out. It’s pretty obvious that mutual defecting gets the biggest jail time and cooperating gets the lightest, at least when we’re talking about both people.** **But if we’re looking for the best strategy for one individual, what we’re really looking for are ways to maximize that person’s benefits while minimizing their maximum cost.**
07:05	**HOST (V.O.)** **For example, let’s pick Mr. Blue. If he cooperates with White, he gets a Reward of a light sentence:**
07:12	BLUE I don’t know anything.
07:14	JUDGE 30 days!
07:16	HOST (V.O.) But if Blue succumbs to the temptation to defect and White cooperates, Blue goes free and White gets the worst punishment -- the Sucker’s Payoff:
07:27	BLUE

	White did it.
07:29	JUDGE 90 days for White! Blue is free to go.
07:32	HOST (V.O.) And if both Blue and White defect, it's the harshest punishment for them both of them.
07:37	BLUE White did it. WHITE Blue did it.
07:40	JUDGE 60 days, the both of you!
07:44	HOST So what's a prisoner to do? If I'm a prisoner, the potential payoffs really define the game. They're ranked in this order - "T" (temptation to defect) is greater than "R" (the reward), which is greater than "P" (the punishment), which is greater than "S" (the sucker's payoff)… And if we plug in values...
08:03	**HOST (V.O.)** **...the payoff matrix clearly shows the stakes and the dilemma, because it seems like choosing to Defect is always the best strategy. In mathematical terms, P is what we call the Minimax solution – a choice that minimizes the maximum loss.**
08:19	HOST Hungarian-American mathematician John Von Neumann described the Minimax solution in 1928 and effectively established the field of game theory.
08:30	**HOST (V.O.)** **Using functional calculus and topology, and chess, Von Neumann proved it possible to work out the best strategy in zero-sum games that would maximize potential gains or minimize potential losses.**
08:45	**HOST (V.O.)** **Von Neumann quickly recognized that his ideas could be applied to the game of business.**
8:50	**HOST (V.O.)** **So in 1944, he teamed up with economist Oskar Morgenstern and** **wrote Theory of Games and Economic Behavior. The book revolutionized the field of economics.**
09:01	**HOST (V.O.)** **At the time, economists focused on how each individual**

	responds to the market and not how individuals interact with each other. Von Neumann and Morgenstern argued that game theory provides a tool to measure how each player's actions influence their rivals.
09:17	HOST With the Minimax Solution, there was at last some mathematical way to help figure out the best strategy in a zero-sum game. But the problem remained - is there a best strategy for a non zero sum game like the Prisoner's Dilemma?
09:31	HOST The complexities of non-zero sum games were of great interest to the mathematician John Nash. In a series of articles published between 1950 and 1953, Nash produced some amazing insights into these kinds of situations.
09:46	**HOST (V.O.)** **While still a student at Princeton Nash realized that in any finite game -- not just a zero-sum game -- there is always a way for players to choose their strategies so that none will wish they had done something else.**
09:59	**HOST (V.O.)** **For the Prisoner's Dilemma, the best strategy is to always Defect, that is, it's a pure D strategy.**
10:05	**HOST (V.O.)** **The Minimax theorem had already showed why in terms of costs and benefits, but Nash's insight was about behavior: if I play my strategy against your strategy, is there a point where changing my strategy won't help me?**
10:19	HOST The answer is yes. Knowing that, and searching for that point, creates what Nash called a strategic equilibrium in the system. And the strategy that creates that equilibrium is now, quite naturally called, the Nash Equilibrium.
10:35	HOST However, this didn't necessarily mean that the payoffs to each individual are desirable... so it still looked like selfish interest was the rule in game theory. But, as we said, people aren't numbers, and they do seem to cooperate, to trust each other, at least sometimes....
11:00	WHITE You ratted me out. BLUE You ratted me out.
11:04	WHITE

	Whatcha readin'? BLUE Book on mathematics.
11:09	WHITE You got a plan for when we get out? BLUE Maybe. WHITE What about that drug store, y'know the one on Broadway? BLUE Didn't we already do that one? WHITE Seems like it's ripe. BLUE I guess. Third time's the charm.
11:29	DAN ROCKMORE: Rock, Paper, Scissors is a game played by children, adults, even prison guards all over the world. But while it's just a game, it's also an interesting mathematical object, and it's the next step in our investigation of game theory.
11:41	ROCKMORE: I'm here with David Krakauer. David is a research professor at the Santa Fe Institute whose work lies at the interface of evolutionary biology, mathematics, and computer science.
11:51	ROCKMORE: So, David, Rock, Paper, Scissors, just a game for prison guards?
11:55	DAVID KRAKAUER Well, no. I mean, what makes this game interesting is there's no best pure strategy solution. If you take Rock, Paper, Scissor, well, let's play it. Let's say I play stone while you play paper. Well, see, paper seems to be better than stone, so I'll play paper. Well, now you play scissor. Well, scissor seems to be the best of all because it's better than the previous move, which is better than the previous move, so it must be the best, but now you play scissor and I play stone, and I win, so you've lost.

	So there's this peculiar property called non-transitivity of the payoff, and that leads to a strange solution where there is no best pure strategy.
12:29	ROCKMORE: There's no best thing for me to do absolutely every time.
12:33	KRAKAUER Exactly. Unconditionally. And so in this game, it turns out the best thing you can do is just play completely randomly. You play each strategy with a probability of 1/3rd.
12:45	ROCKMORE: So I have to randomize. So that randomization is an example of a mixed strategy, is that right?
12:48	KRAKAUER Exactly. Mixed strategy simply refers to the probability of playing any one of the pure strategies. In this case, the pure strategies would be paper, scissor, or rock. And the mixed strategy specifies the probability associated with each pure strategy, so a third, a third, and a third.
13:04	ROCKMORE: Right, and if I deviated from that in any way, then you could exploit that in some fashion.
13:10	KRAKAUER …Yeah, if I saw that, Dan, you liked particularly playing rock, I'd pick up on that cue, and I'd just start playing paper, and then I'd get overall a larger score than you.
13:19	ROCKMORE: Right.
13:20	KRAKAUER …And so we have lots of thoughts in our head and intuitions about things. We're not quite sure which is right, what's superfluous, and what's real. And so mathematics can help to amplify the weak intellectual signal. And so a good example is, you know, what are our intuitions about cooperation? When should we be nice, when should we not cooperate? And using mathematics and computational modeling, Axelrod, University of Michigan, political scientist, in 1978 staged a tournament of computer programs competing in a virtual world over the prisoner's dilemma game.
13:54	ROCKMORE: …so you have a whole collection of people, and everybody's competing against each other, trying to stay out of jail …
13:57	KRAKAUER

	Exactly. So what you have is a large number of computer programs all competing so as to maximize their payoff. And so in the first tournament that was held, 14 computer programs were contributed. And there was one clear winner.
14:13	KRAKAUER And the one that won was tit for tat. And tit for tat just says: "do unto others what they do unto you." And so I just copy your move in the last game.
14:22	ROCKMORE: …so if I cooperated last time…so I'm playing you and if you cooperated last time, then the next game, I'm going to cooperate. If you defected last time, in the next game I'm going to defect when I play you.
14:33	KRAKAUER So here's this hugely complex problem, the problem of cooperation. Somehow you capture the essence of the problem in the prisoner's dilemma matrix, which is this trivial little two-by-two matrix, somehow gets to the heart of the problem. And then you find that the way to do this, to win that game when it's repeated several times, is to play tit for tat and nothing more complex.
14:52	Blue & White It wasn't him, it wasn't him…..
14:56	ROCKMORE: So tit for tat is interesting, but it does seem to have limitations because ultimately, it could also be in one of these anti-cooperative death spirals, if you like: I defect, you defect, I defect.
15:07	Blue & White He did it, he did it….
15:11	KRAKAUER That kind of idiotic solution where you simply copy what the guy did in the last round, it leads to that perpetual defection. And it turns out that when there's some noise or uncertainty, then tit for tat is not the best strategy, so when Axelrod had that tournament, it was working inside a computer, errors were never made. The only uncertainty was what your opponent was going to play. But you always knew exactly what they played once they played it.
15:39	KRAKAUER But let's say that you forgot what they played. So I play you, Dan, and let's say you cooperated, and I think, "did Dan cooperate or defect? I think he defected." So I defect and then you defect. Now it turns out there's an alternative strategy that does better when the world is uncertain, and that strategy --
15:54	ROCKMORE: Which is closer to life, of course.

15:55	KRAKAUER Which is much closer to life, and that strategy is called Pavlov,
15:59	KRAKAUER …named after Pavlov, who did work on conditioning. And specifically on the notion of reinforcement. That if you do something good that's rewarded, you'll do it again. And if you do something bad that's punished, you're less likely to do it again. And so there's a strategy called Pavlov which plays by the so-called "win, stay; lose, shift" rule. And that rule can error correct.
16:20	KRAKAUER So it can take care of this uncertainty.
16:23	KRAKAUER And the intuition there is that if you defect against me, I've lost, so I should shift. And so I shift back to cooperate. And then you see cooperation in the last round, and you cooperate again. And then since -- and then you stay -- you're winning, you stay on that strategy.
16:38	ROCKMORE: So you essentially want to learn from your mistakes. KRAKAUER Exactly. HOST Right
16:41	ROCKMORE: So Nash was actually solving this as a pure math problem, but in fact, it has an evolutionary context, is that right?
16:49	KRAKAUER That's right. So in 1973, an English evolutionary biologist, John Maynard Smith, rediscovered the Nash equilibrium and called it an evolutionary stable strategy. And he was particularly interested in what limits aggression. And it turns out that if you write down a simple game, you can show why it's often the case that more passive, restrained strategies evolve.
17:11	KRAKAUER And the game that he wrote down was called the Hawk/Dove game.
17:14	**HOST (V.O.)** **Imagine we have two populations, one aggressive and one passive.** **Hawks will always fight over a resource, and Doves will not fight under any circumstances. When a Dove meets a Hawk, the Dove always backs down and gives up the resource to the**

	Hawk.
17:34	**HOST (V.O.)** **And when a Hawk fights a Hawk over a resource, the conflict is brutal, the winner takes all, and the loser, well, he ends up injured. The winner gets the reward for this interaction, but because he's suffered a cost in the process, it diminishes the value of that benefit. We can write this out mathematically like this:**
17:56	**HOST (V.O.)** **The benefit of winning the resource, which is B, minus the cost of the fight to get it, which is C. Since a Hawk would win about half the time, the net payoff is B minus C divided by 2.**
18:10	**HOST (V.O.)** **But when a Dove meets a Dove, they share equally with no injury. In other words, they get the benefit half the time but never pay a cost of conflict.**
18:21	***HOST (V.O.)*** **As long as the benefit to be gained from each interaction outweighs the cost of fighting, there is a clear best strategy: be a hawk. But, when the cost of fighting is higher than the benefit to be gained, the logic changes, and doves can succeed. Under these circumstances, the stable population will be a mix of both hawks and doves.**
18:44	ROCKMORE: And do we actually see this in the world in any particular species patterns, things like that?
18:50	KRAKAUER Right. But this is an interesting question, and it relates to how you map highly abstract mathematics of the sort that we're talking about to real-world empirical observations. And I would claim that this kind of mathematics conforms to that model of an intuition amplifier rather than a strategy calculator because it doesn't -- it's so simplified and so abstracted, it tells you why not everyone is mean and aggressive, but it can't tell you precisely how many will be aggressive or non-aggressive.
19:18	ROCKMORE: So this is amazing. So now we have a mathematics that is really beginning to get at the way we think. And that's what we see now in sort of game theory applied to real economics, uncertain payoffs for example, game theory of evolution where you really need inheritance and things like that. So there's still a big world out there for game theory to move into and to change for, so thanks for coming. It's been fascinating.

19:43	KRAKAUER Thank you.
19:45	HOST Game theory can help us understand why animals evolve over time. But it can also help us understand social behavior. Before the 1960's, some scientists thought that the natural selection motto of "survival of the fittest" as applied to behavior would favor the dominance of aggressive behavior...the strong over the weak. …
20:05	HOST Maynard Smith showed that the most evolutionarily stable society is one in which both Hawks and Doves have a role -- which is why natural selection actually works to maintain a balance of different characteristics in a population.
20:21	CRAIG PACKER I'm interesting in discovering why animals behave the way they do, and the only way to do it really is mathematically…
20:30	PACKER My name is Craig Packer. I'm a professor in the Department of Ecology, Evolution and Behavior at the University of Minnesota. Much of my research has been informed by game theory.
20:44	PACKER We're at Wildlife Safari in Winston, Oregon, and we've come to see some lions and see if any of their behaviors illustrate some of these principles of game theory…
20:48	PACKER So the two males are still intact? SARAH Yes, they are.
21:01	PACKER …I started studying lions in the late 1970s on a population of lions that had already been studied for 12 years. Lions are one of the most militantly social species of all mammals. They work together to raise their babies; they often work together to hunt.
21:31	PACKER Our current study area in the Serengeti is 2,000 square kilometers, and we're keeping tabs on 24 different prides of lions. It's actually the most extensive study of any carnivore anywhere in the world.
21:45	PACKER I think evolutionary game theory is a very powerful tool for understanding animal behavior. With animals you have the very simplifying situation that you never can ask them what they're thinking.

	All you can do is rely on the outcomes.
22:01	PACKER Looks like you've got a fairly relaxed group. When is the rut? Sarah: It happened about three weeks ago… PACKER That's what it's all about. I mean the only point of being a male and being so splendid and everything is to get those splendid genes into the next generation.
22:18	PACKER …one of our big questions in studying these lions for the last few decades has been to approach the problem of why it is that lions are the only social cat.
22:30	PACKER And so we're now using a game theoretical approach….
22:33	PACKER …what we're finding is that sociality is much more likely to evolve in a situation where the animals live on very high-quality habitat.
22:45	PACKER They have water, they have food, they have places to hide so you can reach out and grab your prey.
22:50	PACKER What you get then are these singletons now becoming groups defending those territories against anybody else, and that becomes the new ESS, the new evolutionarily stable strategy.
23:04	PACKER When I first started studying lions in the 1970s …
23:08	PACKER …there was always a bias in people that it was a mistake ever to imbue an animal with a complex repertory of behaviors.
23:19	PACKER Maynard Smith with Game Theory comes along and says, "Well if I'm a lion, I live in a world filled with other lions, and so what I get depends on what the other lions are doing."
23:28	PACKER …he brought genetics into the whole story
23:32	PACKER People were convinced that Lions were social because they had to work together to cooperate to catch their prey. And when we did our own research on that subject we found that not only did they not cooperate, but if you thought about it for a few minutes; why should they cooperate? Because every individual in every group no matter how unified the group may appear at first appearance; everybody has their

	own self-interests. And as it happens in a situation like hunting it often is better off if you just notice that ah, my companion or my sister or my mother or whoever is half way to catching that zebra, that looks good, if I just sit still I get a free lunch!
24:19	PACKER …more and more data are showing that animals seem incapable of solving a prisoner's dilemma. They go for the instant gratification. If there's a mutualistic benefit, they always cooperate. If it's not immediately mutualistic, then they don't do it.
24:40	PACKER I study problems. And I love the problems the lions present because they have such a complex social system and they play such a complex role in ecosystems that understanding their behavior is incredibly important. And so always it's the problem that we haven't really addressed yet that's the most exciting.
25:02	HOST So just like with tit-for-tat and Pavlov, the Evolutionary Stable Strategy provides us with a model that, in a sense, buttresses our own intuition about how the world works. Now if we can just keep learning the lessons of Game Theory...
25:18	BOY 1 Hey Mr. Blue? I thought you just got out? WHITE What do you make of it?
25:26	INSPECTOR You guys must be the worst robbers in the city. Just two days out of jail and you are back again. What's your story this time?
25:34	BLUE/WHITE I don't know nothin'. Judge: 30 days
25:44	HOST Game theory forces us to think about choices, strategies and payoffs. Not in a way that reduces us to easily predictable individuals caught in a grid, but in relation to the activity of others.
25:57	HOST In the iterated Prisoner's Dilemma game, it would be great if everybody played a pure cooperate strategy since this would give the greatest payoff. But the temptation to cheat - to buck the system – is there. Maybe that's the point - that math goes beyond our instincts. Our instincts are often wrong and mathematics, carefully considered, can be a guide beyond the "gut." With mathematics we can show that a

	common behavior that we consider foolish, can in fact make considerable sense! Sometimes these "odd" strategies are informally encoded in cultural norms, like the golden rule.
26:35	HOST At its heart, that's perhaps really what game theory is about: the evolution of these rules and norms or institutions that make the best of the difficult situation of living in our world...
26:50	CREDITS